Iterative Ranking from Pair-wise Comparisons

Sahand Negahban, Sewoong Oh, and Devavrat Shah*

Abstract

The question of aggregating pairwise comparisons to obtain a global ranking over a collection of objects has been of interest for a very long time: be it ranking of online gamers (e.g. MSR's TrueSkill system) and chess players, aggregating social opinions, or deciding which product to sell based on transactions. In most settings, in addition to obtaining ranking, finding 'scores' for each object (e.g. player's rating) is of interest for understanding the intensity of the preferences.

In this paper, we propose a novel iterative rank aggregation algorithm for discovering scores for objects (or items) from pairwise comparisons. The algorithm has a natural *random walk* interpretation over the graph of objects with an edge present between a pair of objects if they are compared; the scores turn out to be the stationary probability of this random walk. The algorithm is model independent. To establish the efficacy of our method, however, we consider the popular Bradley-Terry-Luce (BTL) model in which each object has an associated score which determines the probabilistic outcomes of pairwise comparisons between objects. We bound the finite sample error rates between the scores assumed by the BTL model and those estimated by our algorithm. In particular, the number of samples required to learn the score well with high probability depends on the structure of the comparison graph. When the Laplacian of the comparison graph has a strictly positive spectral gap, e.g. each item is compared to a subset of randomly chosen items, this leads to order-optimal dependence on the number of samples. Experimental evaluations on synthetic datasets generated according to the BTL model show that our (model independent) algorithm performs as well as the Maximum Likelihood estimator for that model and outperforms a recently proposed algorithm by Ammar and Shah [AS11].

1 Introduction

Rank aggregation is an important task in a wide range of learning and social contexts arising in recommendation systems, information retrieval, and sports and competitions. Given n items, we wish to infer relevancy scores or an ordering on the items based on partial orderings provided through many (possibly contradictory) samples. Frequently, the available data that is presented to us is in the form of a comparison: player A defeats player B; book A is purchased when books A and B are displayed (a bigger collection of books implies multiple pairwise comparisons); movie A is liked more compared to movie B. From such partial preferences in the form of comparisons, we frequently wish to deduce not only the order of the underlying objects, but also the scores associated with the objects themselves so as to deduce the intensity of the resulting preference order.

For example, the Microsoft TrueSkill engine assigns scores to online gamers based on the outcomes of (pairwise) games between players. Indeed, it assumes that each player has inherent "skill"

*SN and DS are with LIDS and Department of EECS, MIT while SO is with the department of IESE, UIUC. Their email addresses are {sahandn,devavrat}@mit.edu, swoh@illinois.edu.

and the outcomes of the games are used to learn these skill parameters which in turn lead to scores associated with each player. In most such settings, similar model-based approaches are employed.

In this paper, we have set out with the following goal: develop an algorithm for the above stated problem which (a) is computationally simple, (b) works with available (comparison) data only and does not try to fit any model per se, (c) *makes sense* in general, and (d) if the data indeed obeys a reasonable model, then the algorithm should do as well as the best model aware algorithm. The main result of this paper is an affirmative answer to all these questions.

Related work. Most rating based systems rely on users to provide explicit numeric scores for their interests. While these assumptions have led to a flurry of theoretical research for item recommendations based on matrix completion [CR09, KMO10, NW12], it is widely believed that numeric scores provided by individual users are generally inconsistent. Furthermore, in a number of learning contexts as illustrated above, it is simply impractical to ask a user to provide explicit scores.

These observations have led to the need to develop methods that can aggregate such forms of ordering information into relevance ratings. In general, however, designing consistent aggregation methods can be challenging due in part to possible contradictions between individual preferences. For example, if we consider items A, B, and C, one user might prefer A to B, while another prefers B to C, and a third user prefers C to A. Such problems have been well studied as in the work by Condorcet [Con85]. In the celebrated work by Arrow [Arr63], existence of a rank aggregation algorithm with reasonable sets of properties (or axioms) was shown to be impossible.

In this paper, we are interested in a more restrictive setting: we have outcomes of pairwise comparisons between pairs of items, rather than a complete ordering as considered in [Arr63]. Based on those pairwise comparisons, we want to obtain a ranking of items along with a score for each item indicating the intensity of the preference. One reasonable way to think about our setting is to imagine that there is a distribution over orderings or rankings or permutations of items and every time a pair of items is compared, the outcome is generated as per this underlying distribution. With this, our question becomes even harder than the setting considered by Arrow [Arr63] as, in that work, effectively the entire distribution over permutations was already known!

Indeed, such hurdles have not stopped the scientific community as well as practical designers from designing such systems. Chess rating systems and the more recent MSR TrueSkill Ranking system are prime examples. Our work falls precisely into this realm: design algorithms that work well in practice, makes sense in general, and perhaps more importantly, have attractive theoretical properties under common comparative judgment models.

With this philosophy in mind, in recent work, Ammar and Shah [AS11] have presented an algorithm that tries to achieve the goal with which we have set out. However, their algorithm requires information about comparisons between all pairs, and for each pair it requires the exact pairwise comparison 'marginal' with respect to the underlying distribution over permutations. Indeed, in reality, not all pairs of items can typically be compared, and the number of times each pair is compared is also very small. Therefore, while an important step is taken in [AS11], it stops short of achieving the desired goal.

In somewhat related work by Braverman and Mossel [BM08], the authors present an algorithm that produces an ordering based on $O(n \log n)$ pair-wise comparisons on adaptively selected pairs. They assume that there is an underlying true ranking and one observes noisy comparison results. Each time a pair is queried, we are given the true ordering of the pair with probability $1/2 + \gamma$ for some $\gamma > 0$ which does not depend on the items being compared. One limitation of this model is that it does not capture the fact that in many applications, like chess matches, the outcome of a

comparison very much depends on the opponents that are competing.

Such considerations have naturally led to the study of noise models induced by parametric distributions over permutations. An important and landmark model in this class is called the Bradley-Terry-Luce (BTL) model [BT55, Luc59], which is also known as the Multinomial Logit (MNL) model (cf. McFadden [McF73]). It has been the backbone of many practical system designs including pricing in the airline industry [TR05]. Adler, Gemmell, Balter, Karp and Kenyon [AGHB$^+$94] used such models to design adaptive algorithms that select the winner from small number of rounds. Interestingly enough, the (near-)optimal performance of their adaptive algorithm for winner selection is matched by our non-adaptive (model independent) algorithm for assigning scores to obtain global rankings of all players.

Finally, earlier work by Dwork et. al. [DKNS01] propose a number of Markov chain based methods for rank aggregation. The chain named MC3 in that work exactly corresponds to the algorithm presented here. However, our derivation of the algorithm draws from the intuition arising from the recursive update equation discussed in Section 2.2. Furthermore, in this work we provide precise theoretical guarantees predicting the performance of the algorithm.

Our contributions. In this paper, we provide an iterative algorithm that takes the noisy comparison answers between a subset of all possible pairs of items as input and produces scores for each item as the output. The proposed algorithm has a nice intuitive explanation. Consider a graph with nodes/vertices corresponding to the items of interest (e.g. players). Construct a random walk on this graph where at each time, the random walk is likely to go from vertex i to vertex j if items i and j were ever compared; and if so, the likelihood of going from i to j depends on how often i lost to j. That is, the random walk is more likely to move to a neighbor who has more "wins". How frequently this walk visits a particular node in the long run, or equivalently the stationary distribution, is the score of the corresponding item. Thus, effectively this algorithm captures preference of the given item versus all of the others, not just immediate neighbors: the global effect induced by transitivity of comparisons is captured through the stationary distribution.

Such an interpretation of the stationary distribution of a Markov chain or a random walk has been an effective measure of relative importance of a node in wide class of graph problems, popularly known as the *network centrality* [New10]. Notable examples of such network centralities include the random surfer model on the web graph for the version of the PageRank [BP98] which computes the relative importance of a web page, and a model of a random crawler in a peer-to-peer file-sharing network to assign trust value to each peer in EigenTrust [KSGM03].

The computation of the stationary distribution of the Markov chain boils down to 'power iteration' using transition matrix lending to a nice iterative algorithm. Thus, in effect, we have produced an algorithm that (a) is computationally simple and iterative, (b) is model independent and works with the data only, and (c) intuitively makes sense. To establish rigorous properties of the algorithm, we analyze its performance under the BTL model described in Section 2.1.

Formally, we establish the following result: given n items, when comparison results between randomly chosen $O(n\mathsf{poly}(\log n))$ pairs of them are produced as per an (unknown) underlying BTL model, the stationary distribution produced by our algorithm (asymptotically) matches the true score (induced by the BTL model). It should be noted that $\Omega(n \log n)$ is a necessary number of (random) comparisons for any algorithm to even produce a consistent ranking (due to connectivity threshold of random graph). In that sense, up to $\mathsf{poly}(\log n)$ factor, our algorithm is optimal in terms of sample complexity.

In general, the comparisons may not be available between randomly chosen pairs. Let $G =$

$([n], E)$ denote the graph of comparisons between these n objects with an edge $(i, j) \in E$ if and only if objects i and j are compared. In this setting, we establish that with $O(n\mathsf{poly}(\log n)\xi^{-2})$ comparisons, our algorithm learns the true score of the underlying BTL model. Here, ξ is the spectral gap for the Laplacian of G and this how the graph structure of comparisons plays role. Indeed, as a special case when comparisons are chosen at random, the induced graph is Erdös-Rényi for which ξ turns out to be constant, leading to the (order) optimal performance of the algorithm as stated earlier.

To understand the performance of our algorithm compared to the other options, we perform an empirical experimental study. It shows that the performance of our algorithm is identical to the ML estimation of the BTL model. Furthermore, it handsomely outperforms other popular choices including the algorithm by [AS11].

Some remarks about our analytic technique. Our analysis boils down to studying the induced stationary distribution of the random walk or Markov chain corresponding to the algorithm. Like most such scenarios, the only hope to obtain meaningful results for such 'random noisy' Markov chain is to relate it to stationary distribution of a *known* Markov chain. Through recent concentration of measure results for random matrices and comparison technique using Dirichlet forms for characterizing the spectrum of reversible/self-adjoint operators, along with the known expansion property of the random graph, we obtain the eventual result. Indeed, it is the consequence of such powerful results that lead to near-optimal analytic results for random comparison model and characterization of the algorithm's performance for general setting.

The remainder of this paper is organized as follows. In Section 2 we will concretely introduce our model, the problem, and our algorithm. In Section 3 we will discuss our main theoretical results. The proofs will be presented in Section 4.

Notation. In the remainder of this paper, we use C, C', etc. to denote absolute constants, and their value might change from line to line. We use A^T to denote the transpose of a matrix. The Euclidean norm of a vector is denoted by $\|x\| = \sqrt{\sum_i x_i^2}$, and the operator norm of a linear operator is denoted by $\|A\|_2 = \max_x x^T Ax/x^T x$. When we say *with high probability*, we mean that the probability of a sequence of events $\mathcal{A}_n$ goes to one as n grows: $\lim_{n\to\infty} \mathbb{P}(\mathcal{A}_n) = 1$. Also define $[n] = \{1, 2, \ldots, n\}$ to be the set of all integers from 1 to n.

2 Model, Problem Statement, and Algorithm

We now present a concrete exposition of our underlying probabilistic model used in the analysis of our algorithm as well as a formal description of our problem. We then present our explicit random walk approach to ranking.

2.1 Model definition

In this section we discuss our model of comparisons between various items. As alluded to above, for the purpose of establishing analytic properties of the algorithm, we will assume that the outcome of a specific comparison is governed by the BTL model and that we perform the same number of comparisons for each pair that we select. However, the algorithm is itself model-independent

and operates with data generated in arbitrary manner. Namely, the algorithm does not need to estimate model parameters of the BTL model in order to perform the analysis.

Bradley-Terry-Luce model for comparative judgment. When comparing pairs of items from n items of interest, represented as $[n] = \{1, \ldots, n\}$, the Bradley-Terry-Luce model assumes that there is a weight score $w_i \in \mathbb{R}_+$ (i.e. it is a strictly positive real number) associated with each item $i \in [n]$. The outcome of a comparison for pair of items i and j is determined only by the corresponding weights w_i and w_j. Let Y_{ij}^l denote the outcome of the l-th comparison of the pair i and j, such that $Y_{ij}^l = 1$ if j is preferred over i and 0 otherwise. Then, according to the BTL model,

$$Y_{ij}^l = \begin{cases} 1 & \text{with probability } \frac{w_j}{w_i + w_j} \ , \\ 0 & \text{otherwise} \ . \end{cases}$$

Furthermore, conditioned on the score vector $w = \{w_i\}$, it is assumed that the random variables Y_{ij}^l's are independent of one another for all i, j, and l.

Since the BTL model is invariant under the scaling of the scores, an n-dimensional representation of the scores is not unique. Indeed, under the BTL model, a score vector is the equivalence class $[w] = \{w' \in \mathbb{R}^n | w' = a\, w, \text{ for some } a > 0\}$. The outcome of a comparison only depends on the equivalence class of the score vector.

To get a unique representation, we represent each equivalence class by its projection onto the standard orthogonal simplex such that $\sum_i w_i = 1$. This representation naturally defines a distance between two equivalent classes as the Euclidean distance between two projections:

$$d(w, w') \quad \equiv \quad \left\| \frac{1}{\langle w, \mathbb{1} \rangle} w - \frac{1}{\langle w', \mathbb{1} \rangle} w' \right\| .$$

Our main result provides an upper bound on the (normalized) distance between the estimated score vector and the true underlying score vector.

Sampling model. We also assume that we perform a fixed k number of comparisons for all pairs i and j that are considered (e.g. a best of k series). This assumption is mainly to simplify notations, and the analysis as well as the algorithm easily generalizes to the case when we might have a different number of comparisons for different pairs. Given observations of pairwise comparisons among n items according to this sampling model, we define a *comparisons graph* $G = ([n], E, A)$ as a graph of n items where two items are connected if we have comparisons data on that pair and A denotes the weights on each of the edges in E.

2.2 Random walk approach to ranking

In our setting, we will assume that a_{ij} represents the fraction of times object j has been preferred to object i, for example the fraction of times chess player j has defeated player i. Given the notation above we have that $a_{ij} = (1/k) \sum_{l=1}^{k} Y_{ij}^l$. Consider a random walk on a weighted directed graph $G = ([n], E, A)$, where a pair $(i, j) \in E$ if and only if the pair has been compared. The weight edges are defined based on the outcome of the comparisons: $A_{ij} = a_{ij}/(a_{ij} + a_{ji})$ and $A_{ji} = a_{ji}/(a_{ij} + a_{ji})$ (note that $a_{ij} + a_{ji} = 1$ in our setting). We let $A_{ij} = 0$ if the pair has not been compared. Note that by the Strong Law of Large Numbers, as the number $k \to \infty$ the quantity A_{ij} converges to $w_j/(w_i + w_j)$ almost surely.

A random walk can be represented by a time-independent transition matrix P, where $P_{ij} = \mathbb{P}(X_{t+1} = j | X_t = i)$. By definition, the entries of a transition matrix are non-negative and satisfy $\sum_j P_{ij} = 1$. One way to define a valid transition matrix of a random walk on G is to scale all the edge weights by $1/d_{\max}$, where we define $d_{\max}$ as the maximum out-degree of a node. This rescaling ensures that each row-sum is at most one. Finally, to ensure that each row-sum is exactly one, we add a self-loop to each node. Concretely,

$$P_{ij} \quad = \quad \begin{cases} \frac{1}{d_{\max}} A_{ij} & \text{if } i \neq j \ , \\ 1 - \frac{1}{d_{\max}} \sum_{k \neq i} A_{ik} & \text{if } i = j \ . \end{cases} \tag{1}$$

The choice to construct our random walk as above is not arbitrary. In an ideal setting with infinite samples ($k \to \infty$) per comparison the transition matrix P would define a reversible Markov chain under the BTL model. Recall that a Markov chain is reversible if it satisfies the *detailed balance equation*: there exists $v \in \mathbb{R}^n_+$ such that $v_i P_{ij} = v_j P_{ji}$ for all i, j; and in that case, $\pi \in \mathbb{R}^n_+$ defined as $\pi_i = v_i / (\sum_j v_j)$ is its unique stationary distribution. In the ideal setting (say $k \to \infty$), we will have $P_{ij} = \tilde{P}_{ij} \equiv (1/d_{\max}) w_j / (w_i + w_j)$. That is, the random walk will move from state i to state j with probability proportional to the chance that item j is preferred to item i. In such a setting, it is clear that $v = w$ satisfies the reversibility conditions. Therefore, under these ideal conditions it immediately follows that the vector $w / \sum_i w_i$ acts as a valid stationary distribution for the Markov chain defined by $\tilde{P}$, the ideal matrix. Hence, as long as the graph G is connected and at least one node has a self loop then we are guaranteed that our graph has a unique stationary distribution proportional to w. If the Markov chain is reversible then we may apply the spectral analysis of self-adjoint operators, which is crucial in the analysis of the behavior of the method.

In our setting, the matrix P is a noisy version (due to finite sample error) of the ideal matrix $\tilde{P}$ discussed above. Therefore, it naturally suggests the following algorithm as a surrogate. We estimate the probability distribution obtained by applying matrix P repeated starting from any initial condition. Precisely, let $p_t(i) = \mathbb{P}(X_t = i)$ denote the distribution of the random walk at time t with $p_0 = (p_0(i)) \in \mathbb{R}^n_+$ be an arbitrary starting distribution on $[n]$. Then,

$$p_{t+1}^T \quad = \quad p_t^T P \ . \tag{2}$$

In general, the random walk converges to a stationary distribution $\pi = \lim_{t \to \infty} p_t$ which may depend on p_0.When the transition matrix has a unique largest eigenvector (unique stationary distribution), starting from any initial distribution p_0, the limiting distribution π is unique. This stationary distribution π is the top left eigenvector of P, which makes computing π a simple eigenvector computation. Formally, we state the algorithm, which assigns numerical scores to each node, which we shall call *Rank Centrality*:

Rank Centrality
Input: $G = ([n], E, A)$
Output: rank $\{\pi(i)\}_{i \in [n]}$
1: Compute the transition matrix P according to (1);
2: Compute the stationary distribution π (as the limit of (2)).

The stationary distribution of the random walk is a fixed point of the following equation:

$$\pi(i) = \sum_j \pi(j) \frac{A_{ji}}{\sum_\ell A_{i\ell}} \ .$$

This suggests an alternative intuitive justification: an object receives a high rank if it has been preferred to other high ranking objects or if it has been preferred to many objects.

One key question remains: does P have a well defined stationary distribution? Since the Markov chain has a finite state space, there is always a stationary distribution or solution of the above stated fixed-point equations. However, it may not be unique if the Markov chain P is not irreducible. The irreducibility follows easily when the graph is connected and for all edges $(i,j) \in E$, $a_{ij} > 0$, $a_{ji} > 0$. Interestingly enough, we show that the iterative algorithm produces a meaningful solution with near optimal sample complexity as stated in Theorem 2 when the pairs of objects that are compared are chosen at random.

3 Main Results

The main result of this paper derives sufficient conditions under which the proposed iterative algorithm finds a solution that is close to the true solution (under the BTL model) for general model of comparison (i.e. any graph G). This result is stated as Theorem 1 below. In words, the result implies that to learn the true score correctly as per our algorithm, it is sufficient to have number of comparisons scaling as $O(n\mathsf{poly}(\log n)\xi^{-2})$ where ξ is the spectral gap of the Laplacian of the graph G. This result explicitly identifies the role played by the graph structure in the ability of the algorithm to learn the true scores.

In the special case, when the pairs of objects to be compared are chosen at random, that is the induced G is an Erdös-Rényi random graph, the ξ turns out to be constant and hence the resulting number of comparisons required scales as $O(n\mathsf{poly}(\log n))$. This is effectively the optimal sample complexity.

The bounds are presented as the rescaled Euclidean norm between our estimate π and the underlying stationary distribution $\tilde{P}$. This error metric provides us with a means to quantify the relative certainty in guessing if one item is preferred over another. Furthermore, producing such scores are ubiquitous [DMJ10] as they may also be used to calculate the desired rankings. After presenting our main theoretical result we will then provide simulations demonstrating the empirical performance of our algorithm in different contexts.

3.1 Rank Centrality: Error bound for general graphs

Recall that in the general setting, each pair of objects or items are chosen for comparisons as per the comparisons graph $G([n], E)$. For each such pair, we have k comparisons available. The result below characterizes the performance of Rank Centrality for such a general setting.

Before we state the result, we present a few necessary notations. Let d_i denote the degree of node i in G; let the max-degree be denoted by $d_{\max} \equiv \max_i d_i$ and min-degree be denoted by $d_{\min} \equiv \min_i d_i$; let $\kappa \equiv d_{\max}/d_{\min}$. The Laplacian matrix of the graph G is defined as $L = D^{-1}B$ where D is the diagonal matrix with $D_{ii} = d_i$ and B is the adjacency matrix with $B_{ij} = B_{ji} = 1$ if $(i,j) \in E$ and 0 otherwise. The Laplacian, defined thus, can be thought of as a transition matrix of a reversible random walk on graph G: from each node i, jump to one of its neighbors j with equal probability. Given this, it is well known that the Laplacian of the graph has real eigenvalues denoted as

$$-1 \leq \lambda_n(L) \leq \cdots \leq \lambda_1(L) = 1. \tag{3}$$

We shall denote the *spectral gap* of the Laplacian as

$$\xi \equiv 1 - \lambda_{\max}(L) ,$$

where

$$\lambda_{\max}(L) \equiv \max\{\lambda_2(L), -\lambda_n(L)\} . \tag{4}$$

Now we state the result establishing the performance of Rank Centrality.

Theorem 1. *Given n objects and comparison graph $G = ([n], E)$, let each pair $(i, j) \in E$ be compared for k times with outcomes produced as per a BTL model with parameters $w_1, \ldots, w_n$. Then, there exists positive universal constants C and C' such that for $k \geq 4C^2(1 + (b^5\kappa^2/d_{\max}\xi^2)\log n)$, the following bound on the normalized error holds with probability at least $1 - n^{-C'}$:*

$$\frac{\left\|\pi - \tilde{\pi}\right\|}{\|\tilde{\pi}\|} \leq \frac{Cb^{5/2}\kappa}{\xi}\sqrt{\frac{\log n}{k\, d_{\max}}} ,$$

where $\tilde{\pi}(i) = w_i/\sum_\ell w_\ell$ and $b \equiv \max_{i,j} w_i/w_j$. The constant C' can be made as large as desired by increasing the constant C.

3.2 Rank Centrality: Error bound for random graphs

Now we consider the special case when the comparison graph G is an Erdös-Rényi random graph with pair (i, j) being compared with probability d/n. When d is poly-logarithmic in n, we provide a strong performance guarantee. Specifically, the result stated below suggests that with $O(n\mathsf{poly}(\log n))$ comparisons, Rank Centrality manages to learn the true scores with high probability.

Theorem 2. *Given n objects, let the comparison graph $G = ([n], E)$ be generated by selecting each pair (i, j) to be in E with probability d/n independently of everything else. Each such chosen pair of objects is compared k times with the outcomes of comparisons produced as per a BTL model with parameters $w_1, \ldots, w_n$. Then, there exists positive universal constants C, C' and C'' such that when $d \geq C' \log n$, $k \geq C'$, and $k\, d \geq C' b^5 \log n$, the following bound on the error rate holds with probability at least $1 - n^{-C''}$:*

$$\frac{\left\|\pi - \tilde{\pi}\right\|}{\|\tilde{\pi}\|} \leq Cb^{5/2}\sqrt{\frac{\log n}{k\, d}} ,$$

where $\tilde{\pi}(i) = w_i/\sum_\ell w_\ell$ and $b \equiv \max_{i,j} w_i/w_j$. The C'' can be made as large as desired by increasing the constants C and C'.

3.3 Remarks

Some remarks are in order. First, Theorem 2 implies that as long as we choose $d = \Theta(\log^2 n)$ and $k = \omega(1)$ the error goes to 0. For $k = \Theta(\log n)$, it goes down at a rate $1/\log n$ as n increases. Since we are sampling each of the $\binom{n}{2}$ pairs with probability d/n and then obtaining k comparisons per pair, we obtain $O(n \log^3 n)$ comparisons in total with $k = \Theta(\log n)$ and $d = \Theta(\log^2 n)$. Due to

classical results on Erdos-Renyi graphs, the induced graph G is connected with high probability only when total number of pairs sampled scales as $\Omega(n \log n)$–we need at least those many comparisons. Thus, our result can be sub-optimal only up to $\log^2 n$ ($\log^\epsilon n$ if $k = \log^\epsilon n$ and $d = \log n$).

Second, the b parameter should be treated as constant. It is the *dynamic* range in which we are trying to resolve the uncertainty between scores. If b were scaling with n, then it would be really easy to differentiate scores of items that are at the two opposite end of the dynamic range; in which case one could focus on differentiating scores of items that have their parameter values near-by. Therefore, the interesting and challenging regime is where b is constant and not scaling.

Third, for a general graph, Theorem 1 implies that by choice of $kd_{\max} = O(\kappa^2 \xi^{-2} \log n)$, the true scores can be learnt by Rank Centrality. That is, effectively the Rank Centrality algorithm requires $O(n\kappa^2\xi^{-2}\mathsf{poly}(\log n))$ comparisons to learn scores well. Ignoring κ, the graph structure plays a role through ξ^{-2}, the squared inverse of the spectral gap of Laplacian of G, in dictating the performance of Rank Centrality. A reversible natural random walk on G, whose transition matrix is the Laplacian, has its mixing time scaling as ξ^{-2} (precisely, relaxation time). In that sense, the mixing time of natural random walk on G ends up playing an important role in the ability of Rank Centrality to learn the true scores.

3.4 Experimental Results

Under the BTL model, define an error metric of an estimated ordering σ as the weighted sum of pairs (i, j) whose ordering is incorrect:

$$D_w(\sigma) \;=\; \Big\{ \frac{1}{2n\|w\|^2} \sum_{i<j} (w_i - w_j)^2 \, \mathbb{I}\big((w_i - w_j)(\sigma_i - \sigma_j) > 0\big) \Big\}^{1/2},$$

where $\mathbb{I}(\cdot)$ is an indicator function. This is a more natural error metric compared to the Kemeny distance, which is an unweighted version of the above sum, since $D_w(\cdot)$ is less sensitive to errors between pairs with similar weights. Further, assuming without loss of generality that w is normalized such that $\sum_i w_i = 1$, the next lemma connects the error in $D_w(\cdot)$ to the bound provided in Theorem 2. Hence, the same upper bound holds for D_w error. A proof of this lemma is provided in Section 4.4.

Lemma 3.1. *Let σ be an ordering of n items induced by a scoring π. Then, $D_w(\sigma) \le \|w - \pi\|/\|w\|$.*

For a fixed $n = 400$ and a fixed $b = 10$, Figure. 1 illustrates how the error scales with two problem parameters: varying the number of comparisons per pair with fixed $d = 10 \log n$ (left) and varying the sampling probability with fixed $k = 32$ (right). The ML estimator directly maximizes the likelihood assuming the BTL model [LRF57]. If we reparameterize the problem so that $\theta_i = \log(w_i)$ then we obtain our estimates $\widehat{\theta}$ by solving the convex program

$$\widehat{\theta} \in \arg\min_{\theta} \sum_{(i,j) \in E} \sum_{l=1}^{k} \log(1 + \exp(\theta_j - \theta_i)) - Y_{ij}^l (\theta_j - \theta_i),$$

which is pair-wise logistic regression. This choice is optimal in the asymptotic setting, however for fixed-samples there do not exist theoretical guarantees for recovering the transformed scores θ_i. The method Count Wins, proposed recently by [AS11], scores an item by counting the number of

wins divided by the total number of comparisons. Ratio Matrix assigns scores according to the top eigenvector of a matrix, whose (i, j)-th entry is a_{ij}/a_{ji} [Saa03]. As we see in Figure 1, the error achieved by our Random Walk approach is comparable to that of ML estimator, and vanishes at the rate of $1/\sqrt{k}$ as predicted by our main result. Interestingly, for fixed d, both the Count Wins and Ratio Matrix algorithms have strictly positive error even if we take $k \to \infty$. The figure on the right illustrates that the error scales as $1/\sqrt{d}$ as expected from our main result.

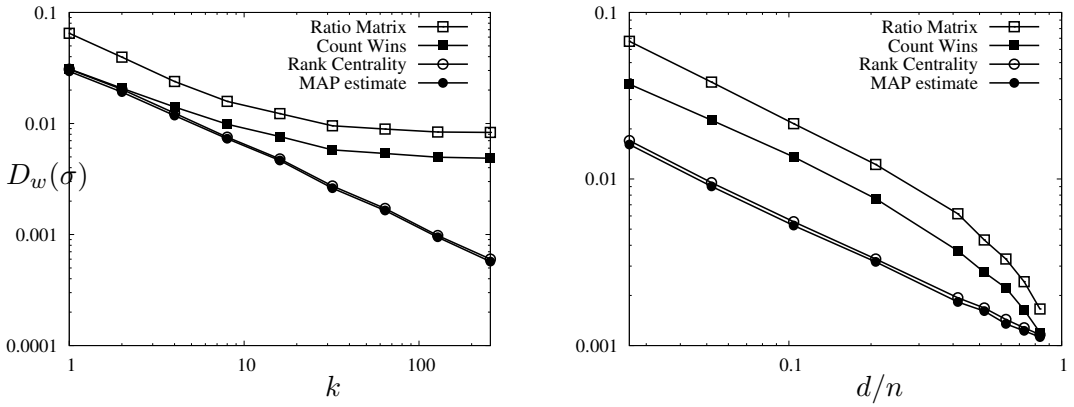

Figure 1: Average error $D_w(\sigma)$ of orderings from four rank aggregation algorithms, averaged over 20 instances. In the figure on the right we assume that d and n are fixed while we increase k. The figure on the right takes $k = 32$ fixed and lets d increase.

To test our algorithm on real data, we used a public dataset collected from an online polling on Washington Post[1] from December 2010 to January 2011. Using allourideas[2] platform developed by [SL12], they asked who had the worst year in Washington, where each user was asked to compare a series of randomly selected pairs of political entities. There are 67 political entities in the dataset, and the resulting graph is a complete graph on these 67 nodes. We used Rank Centrality and Count Wins to aggregate this data. Since we do not have the ground truth for real datasets, we used the ranking we get on the full complete graph as the ground truth. This gives two different ground truth rankings for each algorithm. This ground truth is compared to a ranking we get from only a subset of the data, which is generated by sampling each edge with a given sampling rate and revealing only the data on those sampled edges. We want to measure how much each algorithm is effected by eliminating edges from the complete graph. Let σ_{GT} be the ranking we get by applying our choice of rank aggregation algorithm to the complete dataset, and σ_{Sample} be the ranking we get from sampled dataset. To measure the resulting error in the ranking, we use the following metric:

$$D_{L_1}(\sigma_{\mathrm{GT}}, \sigma_{\mathrm{Sample}}) \quad = \quad \frac{1}{n} \sum_i |\sigma_{\mathrm{GT}}(i) - \sigma_{\mathrm{Sample}}(i)| \; .$$

Figure 2 illustrates that, compared to Count Wins, Rank Centrality is less sensitive to sampling the dataset, and hence more robust when available comparisons data is limited.

[1] `http://www.washingtonpost.com/wp-srv/interactivity/worst-year-voting.html`
[2] `http://www.allourideas.org`

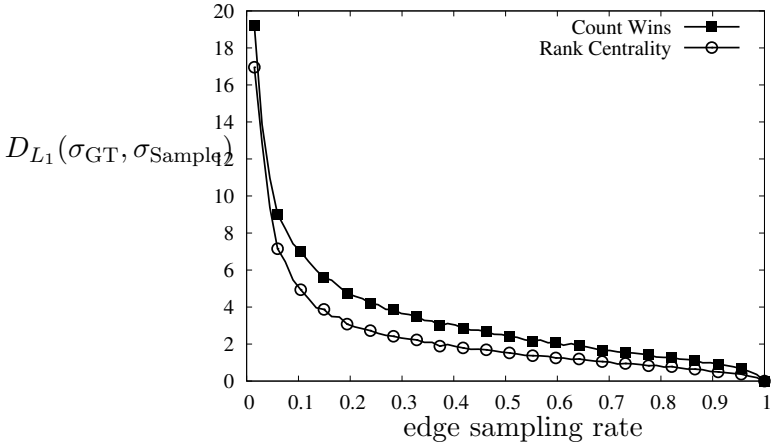

Figure 2: Experimental result on a real dataset shows that Rank Centrality is less sensitive to having limited data.

3.5 Information-theoretic lower bound

In previous sections, we presented the achievable error rate based on a particular low-complexity algorithm. In this section, we ask how this bound compares to the fundamental limit under BTL model.

Our result in Theorem 2 provides an upper bound on the achievable error rate between estimated scores and the true underlying scores. We provide a constructive argument to lower bound the minimax error rate over a class of BTL models. Concretely, we consider the scores coming from a simplex with bounded dynamic range defined as

$$\mathcal{S}_b \quad \equiv \quad \left\{ \tilde{\pi} \in \mathbb{R}^n \;\middle|\; \sum_{i \in [n]} \tilde{\pi}_i = 1 \;,\; \max_{i,j} \frac{\tilde{\pi}_i}{\tilde{\pi}_j} \leq b \right\} .$$

We constrain the scores to be on the simplex, because we represent the scores by its projection onto the standard simplex as explained in Section 2.1. Then, we can prove the following lower bound on the minimax error rate.

Theorem 3. *Consider a minimax scenario where we first choose an estimator π that estimates the BTL weights from given observations and, for this particular estimator π, nature chooses the worst-case true BTL weights $\tilde{\pi}$. Then, we can show that for any estimator π that we choose, there exists a true score vector $\tilde{\pi}$ with dynamic range at most b such that no algorithm can achieve an expected normalized error smaller than the following minimax lower bound:*

$$\inf_{\pi} \sup_{\tilde{\pi} \in \mathcal{S}_b} \frac{\mathbb{E}\big[\, \|\pi - \tilde{\pi}\| \,\big]}{\|\tilde{\pi}\|} \quad \geq \quad \frac{b-1}{240\sqrt{10}(b+1)} \frac{1}{\sqrt{kd}} \;, \tag{5}$$

where the infimum ranges over all estimators π that are measurable functions over the observations, we observe the outcomes of k comparisons for each pair of items, and we compare each pair of items with probability d/n.

By definition the dynamic range is always at least one. When $b = 1$, we can trivially achieve a minimax rate of zero. Since the infimum ranges over all measurable functions, it includes a trivial estimator which always outputs $(1/n)\mathbb{1}$ regardless of the observations, and this estimator achieves zero error when $b = 1$. In the regime where the dynamic range b is bounded away from one and bounded above by a constant, Theorem 3 establishes that the upper bound obtained in Theorem 2 is minimax-optimal up to factors logarithmic in the number of items n.

4 Proofs

We may now present proofs of Theorems 1 and 2. We first present a proof of convergence for general graphs in Theorem 1. This result follows from a lemma that we state below, which shows that our algorithm enjoys convergence properties that result in useful upper bounds. The lemma is made general and uses standard techniques of spectral theory. The main difficulty arises in establishing that the Markov chain P satisfies certain properties that we will discuss below. In order to show that these properties hold we must rely on the specific model that allows us to ultimately establish error bounds that hold with high probability. Given the proof for the general graph, Theorem 2 follows by showing that in the case of Erdös-Renyi graphs, the necessary conditions are satisfied with high probability. Then, we provide a proof of the information-theoretic lower bound.

4.1 Algorithm convergence for general graphs

In this section, we characterize the error rate achieved by our ranking algorithm. Given the random Markov chain P, where the randomness comes from the outcome of the comparisons, we will show that it does not deviate too much from its expectation $\tilde{P}$, where we recall is defined as

$$\tilde{P}_{ij} \quad = \quad \begin{cases} \frac{1}{d_{\max}} \frac{w_j}{w_i + w_j} & \text{if } i \neq j \ , \\ 1 - \frac{1}{d_{\max}} \sum_{\ell \neq i} \frac{w_\ell}{w_i + w_\ell} & \text{if } i = j \end{cases}$$

for all $(i, j) \in E$ and $\tilde{P}_{ij} = 0$ otherwise.

Recall from the discussion following equation (1) that the transition matrix P used in our ranking algorithm has been carefully chosen such that the corresponding expected transition matrix $\tilde{P}$ has two important properties. First, the stationary distribution of $\tilde{P}$, which we denote with $\tilde{\pi}$ is proportional to the weight vectors w. Furthermore, when the graph is connected and has self loops (which at least one exists), this Markov chain is irreducible and aperiodic so that the stationary distribution is unique. The next important property of $\tilde{P}$ is that it is reversible–$\tilde{\pi}(i)\tilde{P}_{ij} = \tilde{\pi}(j)\tilde{P}_{ji}$. This observation implies that the operator $\tilde{P}$ is symmetric in an appropriate defined inner product space. The symmetry of the operator $\tilde{P}$ will be crucial in applying ideas from spectral analysis to prove our main results.

Let Δ denote the fluctuation of the transition matrix around its mean, such that $\Delta \equiv P - \tilde{P}$. The following lemma bounds the deviation of the Markov chain after t steps in terms of two important quantities: the spectral radius of the fluctuation $\|\Delta\|_2$ and the spectral gap $1 - \lambda_{\max}(\tilde{P})$, where

$$\lambda_{\max}(\tilde{P}) \quad \equiv \quad \max\{\lambda_2(\tilde{P}), -\lambda_n(\tilde{P})\} \ .$$

Since $\lambda(\tilde{P})$'s are sorted, $\lambda_{\max}(\tilde{P})$ is the second largest eigenvalue in absolute value.

Lemma 4.1. *For any Markov chain $P = \tilde{P} + \Delta$ with a reversible Markov chain $\tilde{P}$, let p_t be the distribution of the Markov chain P when started with initial distribution p_0. Then,*

$$\frac{\left\|p_t - \tilde{\pi}\right\|}{\|\tilde{\pi}\|} \quad \leq \quad \rho^t \frac{\|p_0 - \tilde{\pi}\|}{\|\tilde{\pi}\|} \sqrt{\frac{\tilde{\pi}_{\max}}{\tilde{\pi}_{\min}}} + \frac{1}{1-\rho} \|\Delta\|_2 \sqrt{\frac{\tilde{\pi}_{\max}}{\tilde{\pi}_{\min}}} \,. \tag{6}$$

where $\tilde{\pi}$ is the stationary distribution of $\tilde{P}$, $\tilde{\pi}_{\min} = \min_i \tilde{\pi}(i)$, $\tilde{\pi}_{\max} = \max_i \tilde{\pi}(i)$, and $\rho = \lambda_{\max}(\tilde{P}) + \|\Delta\|_2 \sqrt{\tilde{\pi}_{\max}/\tilde{\pi}_{\min}}$.

The above result provides a general mechanism for establishing error bounds between an estimated stationary distribution π and the desired stationary distribution $\tilde{\pi}$. It is worth noting that the result only requires control on the quantities $\|\Delta\|_2$ and $1 - \rho$. We may now state two technical lemmas that provide control on the quantities $\|\Delta\|_2$ and $1 - \rho$, respectively.

Lemma 4.2. *For $k \geq 13$ and $kd_{\max} \geq C \log n$ with appropriately large constant C, the error matrix $\Delta = P - \tilde{P}$ satisfies*

$$\|\Delta\|_2 \quad \leq \quad C' \sqrt{\frac{\log n}{k \, d_{\max}}}$$

with probability at least $1 - n^{-C''}$: constant C'' can be made large at the cost of possibly making C and C' larger.

The next lemma provides our desired bound on $1 - \rho$.

Lemma 4.3. *When $\|\Delta\|_2 \leq C\sqrt{\log n/(kd_{\max})}$ and $k \geq 4C^2 b^5 d_{\max} \log n (1/d_{\min}\xi)^2$, the spectral radius satisfies*

$$1 - \rho \quad \geq \quad \frac{\xi d_{\min}}{b^2 d_{\max}} \,.$$

With the above results in hand we may now proceed with the proof of Theorem 1.

When there is a positive spectral gap such that $\rho < 1$, the first term in (6) vanishes as t grows. The rest of the first term is bounded and independent of t. Formally, we have

$$\tilde{\pi}_{\max}/\tilde{\pi}_{\min} \leq b \,, \quad \|\tilde{\pi}\| \geq 1/\sqrt{n} \,, \quad \text{and } \|p_0 - \tilde{\pi}\| \leq 2 \,,$$

by the assumption that $\max_{i,j} w_i/w_j \leq b$ and the fact that $\tilde{\pi}(i) = w_i/(\sum_j w_j)$. Hence, the error between the distribution at the t^{th} iteration p^t and the true stationary distribution $\tilde{\pi}$ is dominated by the second term in equation (6). Substituting the bounds in Lemma 4.2 and Lemma 4.3, the dominant second term in equation (6) is bounded by

$$\lim_{t \to \infty} \frac{\left\|p_t - \tilde{\pi}\right\|}{\|\tilde{\pi}\|} \quad \leq \quad \frac{C \, b^{5/2}}{\xi d_{\min}} \sqrt{\frac{d_{\max} \log n}{k}} \,.$$

In fact, we only need $t = \Omega(\log n + \log b + \log(d_{\max} \log n/(d_{\min}^2 k \xi^2)))$ to ensure that the above bound holds up to a constant factor. This finishes the proof of Theorem 1. Notice that in order for this result to hold, we need the following two conditions: $kd_{\max} \geq C' \log n$ for Lemma 4.2 and $k \geq 4C^2 b^5 d_{\max} \log n (1/d_{\min}\xi)^2$ for Lemma 4.3. Since $b \geq 1$, $d_{\max} \geq d_{\min}$, and $0 \geq \xi \leq 1$, the second condition always implies the first for any choice of $4C^2 \geq C'$.

4.1.1 Proof of Lemma 4.1

Due to the reversibility of $\tilde{P}$, we can view it as a self-adjoint operator on an appropriately defined inner product space. This observation allows us to apply the well-understood spectral analysis of self-adjoint operators. In order to establish this fact define an inner product space $L^2(\tilde{\pi})$ as a space of n-dimensional vectors with

$$\langle a, b\rangle_{\tilde{\pi}} = \sum_{i=1}^{n} a_i \tilde{\pi}_i b_i \ .$$

Similarly, we define $\|a\|_{\tilde{\pi}} = \sqrt{\langle a, a\rangle_{\tilde{\pi}}}$ as the 2-norm in $L^2(\tilde{\pi})$. For a self-adjoint operator A in $L^2(\tilde{\pi})$, we define $\|A\|_{\tilde{\pi},2} = \max_a \|Aa\|_{\tilde{\pi}}/\|a\|_{\tilde{\pi}}$ as the operator norm. These norms are related to the corresponding norms in the Euclidean space through the following inequalities.

$$\sqrt{\tilde{\pi}_{\min}}\,\|a\| \ \leq \ \|a\|_{\tilde{\pi}} \ \leq \ \sqrt{\tilde{\pi}_{\max}}\,\|a\| \ , \tag{7}$$

$$\sqrt{\frac{\tilde{\pi}_{\min}}{\tilde{\pi}_{\max}}}\,\|A\|_2 \ \leq \ \|A\|_{\tilde{\pi},2} \ \leq \ \sqrt{\frac{\tilde{\pi}_{\max}}{\tilde{\pi}_{\min}}}\,\|A\|_2 \ . \tag{8}$$

A reversible Markov chain $\tilde{P}$ is self-adjoint in $L^2(\tilde{\pi})$. To see this, define a closely related symmetric matrix $S = \tilde{\Pi}^{1/2}\tilde{P}\tilde{\Pi}^{-1/2}$, where $\tilde{\Pi}$ is a diagonal matrix with $\tilde{\Pi}_{ii} = \tilde{\pi}(i)$. The assumption that $\tilde{P}$ is reversible, i.e. $\tilde{\pi}(i)\tilde{P}_{ij} = \tilde{\pi}(j)\tilde{P}_{ji}$, implies that S is symmetric, and it follows that $\tilde{P}$ is self-adjoint in $L^2(\tilde{\pi})$.

Further, the asymmetric matrix $\tilde{P}$ and the symmetric matrix S have the same set of eigenvalues. By Perron-Frobenius theorem, the eigenvalues are at most one. Let $1 = \lambda_1 \geq \lambda_2 \geq \ldots \geq \lambda_n \geq -1$ be the eigenvalues, and let u_i be the left eigenvector of S corresponding to λ_i. Then the ith left eigenvector of $\tilde{P}$ is $v_i = \tilde{\Pi}^{1/2}u_i$. Since the first left eigenvector of $\tilde{P}$ is the stationary distribution, i.e. $v_1 = \tilde{\pi}$, we have $u_1(i) = \tilde{\pi}(i)^{1/2}$.

For the Markov chain $P = \tilde{P} + \Delta$, where $\tilde{P}$ is a reversible Markov chain such that $\tilde{\pi}^T\tilde{P} = \tilde{\pi}$, we let $p_t^T = p_{t-1}^T P$. Then,

$$p_t^T - \tilde{\pi}^T \ = \ (p_{t-1} - \tilde{\pi})^T(\tilde{P} + \Delta) + \tilde{\pi}^T\Delta \ . \tag{9}$$

Define $S_0 = \lambda_1 u_1 u_1^T$ to be the rank-1 projection of S, and a corresponding matrix $\tilde{P}_0 = \tilde{\Pi}^{-1/2}S_0\tilde{\Pi}^{1/2}$. Using the fact that $(p_\ell - \tilde{\pi})^T\tilde{\Pi}^{-1/2}u_1 = (p_\ell - \tilde{\pi})^T\mathbb{1} = 0$ for any probability distribution p_ℓ, we get $(p_\ell - \tilde{\pi})^T\tilde{P}_0 = (p_\ell - \tilde{\pi})^T\tilde{\Pi}^{-1/2}u_1\lambda_1 u_1^T\tilde{\Pi}^{1/2} = 0$. Then, from (9) we get

$$p_t^T - \tilde{\pi}^T \ = \ (p_{t-1} - \tilde{\pi})^T(\tilde{P} - \tilde{P}_0 + \Delta) + \tilde{\pi}^T\Delta \ .$$

By definition of $\tilde{P}_0$, it follows that $\|\tilde{P} - \tilde{P}_0\|_{\tilde{\pi},2} = \|S - S_0\|_2 = \lambda_{\max}$. Let $\rho = \lambda_{\max} + \|\Delta\|_{\tilde{\pi},2}$, then

$$\begin{aligned}
\|p_t - \tilde{\pi}\|_{\tilde{\pi}} \ &\leq \ \|p_{t-1} - \tilde{\pi}\|_{\tilde{\pi}}(\|\tilde{P} - \tilde{P}_0\|_{\tilde{\pi},2} + \|\Delta\|_{\tilde{\pi},2}) + \|\tilde{\pi}^T\Delta\|_{\tilde{\pi}} \\
&\leq \ \rho^t\|p_0 - \tilde{\pi}\|_{\tilde{\pi}} + \sum_{\ell=0}^{t-1}\rho^{t-1-\ell}\|\tilde{\pi}^T\Delta\|_{\tilde{\pi}} \ .
\end{aligned}$$

Dividing each side by $\|\tilde{\pi}\|$ and applying the bounds in (7) and (8), we get

$$\frac{\|p_t - \tilde{\pi}\|}{\|\tilde{\pi}\|} \ \leq \ \rho^t\sqrt{\frac{\tilde{\pi}_{\max}}{\tilde{\pi}_{\min}}}\frac{\|p_0 - \tilde{\pi}\|}{\|\tilde{\pi}\|} + \sum_{\ell=0}^{t-1}\rho^{t-1-\ell}\sqrt{\frac{\tilde{\pi}_{\max}}{\tilde{\pi}_{\min}}}\frac{\|\tilde{\pi}^T\Delta\|}{\|\tilde{\pi}\|} \ .$$

This finishes the proof of the desired claim.

4.1.2 Proof of Lemma 4.2

Our interest is in bounding $\|\Delta\|_2$. Now $\Delta = P - \tilde{P}$ so that for $1 \le i, j \le n$,

$$\Delta_{ij} = \frac{1}{kd_{\max}} C_{ij}, \tag{10}$$

where C_{ij} is distributed as per $B(k, p_{ij}) - kp_{ij}$ if $(i,j) \in E$ and $C_{ij} = 0$ otherwise. Here $B(k, p_{ij})$ is a Binomial random variable with parameter k and $p_{ij} \equiv \frac{w_j}{w_i + w_j}$. It should be noted that $C_{ij} + C_{ji} = 0$ and C_{ij} are independent across all the pairs with $i < j$. For $1 \le i \le n$

$$\begin{aligned} \Delta_{ii} = P_{ii} - \tilde{P}_{ii} \; &= \; \big(1 - \sum_{j \neq i} P_{ij}\big) - \big(1 - \sum_{j \neq i} \tilde{P}_{ij}\big) \; = \sum_{j \neq i} \tilde{P}_{ij} - P_{ij} \\ &= - \sum_{j \neq i} \Delta_{ij}. \end{aligned} \tag{11}$$

Given the above dependence between diagonal and off-diagonal entries, we shall bound $\|\Delta\|_2$ as follows: let D be the diagonal matrix with $D_{ii} = \Delta_{ii}$ for $1 \le i \le n$ and $\bar{\Delta} = \Delta - D$. Then,

$$\|\Delta\|_2 = \|D + \bar{\Delta}\|_2 \; \le \|D\|_2 + \|\bar{\Delta}\|_2. \tag{12}$$

We shall establish the bound of $O\big(\sqrt{\frac{\log n}{kd_{\max}}}\big)$ for both $\|D\|_2$ and $\|\bar{\Delta}\|_2$ to establish the Lemma 4.2.

Bounding $\|D\|_2$. Since D is a diagonal matrix, $\|D\|_2 = \max_i |D_{ii}| = \max_i |\Delta_{ii}|$. For a given fixed i, as per (10)-(11), $kd_{\max}\Delta_{ii}$ can be expressed as summation of at most $kd_{\max}$ independent, zero-mean random variables taking values in the range of at most 1. Therefore, by an application of Azuma-Hoeffding's inequality, it follows that

$$\mathbb{P}\big(kd_{\max}|\Delta_{ii}| > t\big) \le 2 \exp\big(-\frac{t^2}{2kd_{\max}}\big). \tag{13}$$

By selection of $t = C\sqrt{kd_{\max} \log n}$ for appropriately large constant, it follows from above display that

$$\mathbb{P}\Big(|\Delta_{ii}| > C\sqrt{\frac{\log n}{kd_{\max}}}\Big) \le O\big(n^{-2C}\big). \tag{14}$$

In summary, we have

$$\|D\|_2 \le C\sqrt{\frac{\log n}{kd_{\max}}}. \tag{15}$$

Bounding $\|\bar{\Delta}\|_2$ *when* $d_{\max} \le \log n$. Towards this goal, we shall make use of the following standard inequality: for any square matrix M,

$$\|M\|_2 \le \sqrt{\|M\|_1 \|M\|_\infty}, \tag{16}$$

where $\|M\|_1 = \max_i \sum_j |M_{ij}|$ and $\|M\|_\infty = \|M^T\|_1$. In words, $\|M\|_2^2$ is bounded above by product of the maximal row-sum and column-sum of absolute values of M. Since Δ_{ij} and Δ_{ji} are identically

distributed and entries along each row (and hence each column) are independent, it is sufficient to obtain a high probability bound ($\geq 1 - 1/\mathsf{poly}(n)$) for maximal row-sum of absolute values of $\bar{\Delta}$; exactly the same bound will apply for column-sum using a union bound.

To that end, consider the sum of the absolute values of the ith row-sum of $\bar{\Delta}$ and for simplicity let us denote it by R_i. Then,

$$R_i = \frac{1}{kd_{\max}} \sum_{j \neq i} |C_{ij}|, \tag{17}$$

where recall that $C_{ij} = X_{ij} - kp_{ij}$ with X_{ij} an independent Binomial random variable with parameters k, p_{ij}. Therefore, for any $s > 0$,

$$\begin{aligned}
\mathbb{P}\big(R_i > s\big) &= \mathbb{P}\big(\sum_j |C_{ij}| > kd_{\max} s\big) \\
&\leq \mathbb{E}[\exp(\sum_j \theta |C_{ij}|)] \exp(-\theta k d_{\max} s), \quad \text{for any } \theta > 0, \\
&= \exp(-\theta k d_{\max} s) \prod_j \mathbb{E}[\exp(\theta |C_{ij}|)].
\end{aligned} \tag{18}$$

Next, we bound $\mathbb{E}[\exp(\theta|C_{ij}|)]$. To that end, observe that for any $x \in \mathbb{R}$ and $\theta > 0$,

$$\exp(\theta|x|) \leq \exp(\theta x) + \exp(-\theta x).$$

From this, it follows that

$$\mathbb{E}[\exp(\theta|C_{ij}|)] \leq \mathbb{E}[\exp(\theta C_{ij})] + \mathbb{E}[\exp(-\theta C_{ij})]. \tag{19}$$

Now for any $\phi \in \mathbb{R}$, using the fact that X_{ij} is Binomial distribution and $1 + x \leq \exp(x)$ for any $x \in \mathbb{R}$, we have

$$\begin{aligned}
\mathbb{E}[\exp(\phi C_{ij})] &= \exp(-\phi k p_{ij}) \big(1 + p_{ij}(\exp(\phi) - 1)\big)^k \\
&\leq \exp(-\phi k p_{ij}) \exp\big(k p_{ij}(\exp(\phi) - 1)\big).
\end{aligned} \tag{20}$$

Using second-order Taylor's expansion, for any $\phi \in [-\ln 4/3, \ln 4/3]$, we obtain that

$$|\exp(\phi) - 1 - \phi| \leq \frac{2}{3}\phi^2. \tag{21}$$

Using above display in (20), we have ave

$$\mathbb{E}[\exp(\phi C_{ij})] \leq \exp\big(2kp_{ij}\phi^2/3\big). \tag{22}$$

From (19) and (22), we have that for $0 \leq \theta \leq \ln 4/3$,

$$\mathbb{E}[\exp(\theta|C_{ij}|)] \leq 2\exp\big(2k\theta^2/3\big). \tag{23}$$

Replacing (23) in (18) and recalling the fact the degree of node $d_i \leq d_{\max}$, we have that for $0 \leq \theta \leq \ln 4/3$,

$$\mathbb{P}\big(R_i > s\big) \leq \exp\big(-\theta k d_{\max} s + 2k d_{\max}\theta^2/3 + d_{\max}\ln 2\big). \tag{24}$$

Using (24), the optimal choice of θ is $\theta = (3/4)s$. Choosing $s = \sqrt{\frac{8}{3kd_{\max}}\big(c\ln n + d_{\max}\ln 2\big)}$, for a given $c > 1$, we obtain

$$\mathbb{P}\Big(R_i > \sqrt{\frac{8}{3kd_{\max}}\big(c\ln n + d_{\max}\ln 2\big)}\Big) \leq n^{-c}. \tag{25}$$

To ensure that $\theta = (3/4)s \leq \ln 4/3$, we need

$$\sqrt{\frac{8}{3kd_{\max}}\big(c\ln n + d_{\max}\ln 2\big)} \leq \frac{4}{3}\ln\frac{4}{3}, \tag{26}$$

which holds when $kd_{\max} \geq C\log n$ and $k \geq 13$ for some positive constant C. From above, and an application of union bound across rows and columns, it follows that with probability at least $1 - O(n^{-c+1})$, as long as $kd_{\max} = \Omega(\log n)$ with appropriately large enough constant, we have that

$$\|\bar{\Delta}\|_2 \leq c'\sqrt{\frac{\log n + d_{\max}}{kd_{\max}}}, \tag{27}$$

for an appropriate choice of constant c'. Note that the above inequality reduces to the desired claim of Lemma 4.2 for any $d_{\max} = O(\log n)$.

Bounding $\|\bar{\Delta}\|_2$ when $d_{\max} \geq \log n$. Towards this goal, we shall make use of the recent results on the concentration of sum of independent random matrices. For completeness, we recall the following result [Tro11].

Lemma 4.4 (Theorem 6.2 [Tro11])**.** *Consider a finite sequence $\{\tilde{Z}^{ij}\}_{i<j}$ of independent random self-adjoint matrices with dimensions $n \times n$. Assume that*

$$\mathbb{E}[\tilde{Z}^{ij}] = 0 \quad \text{and} \quad \mathbb{E}(\tilde{Z}^{ij})^p \preceq \frac{p!}{2}R^{p-2}(\tilde{A}^{ij})^2 \quad , \text{ for } p = 2, 3, 4, \ldots$$

Define $\tilde{\sigma}^2 \equiv \|\sum_{i>j}(\tilde{A}^{ij})^2\|_2$. Then, for all $t \geq 0$,

$$\mathbb{P}\Big(\Big\|\sum_{i<j}\tilde{Z}^{ij}\Big\|_2 \geq t\Big) \;\leq\; 2n\exp\Big\{\frac{-t^2/2}{\tilde{\sigma}^2 + Rt}\Big\}.$$

We wish to prove concentration results on $\bar{\Delta} = \Delta - D = \sum_{i<j} Z^{ij}$ where

$$Z^{ij} \;=\; (e_ie_j^T - e_je_i^T)(P_{ij} - \tilde{P}_{ij}) \quad \text{for } (i,j) \in E\,,$$

and $Z^{ij} = 0$ if i and j are not connected. The Z^{ij}'s as defined are zero-mean and independent, however, they are not self-adjoint. Nevertheless, we can symmetrize it by applying the dilation ideas presented in the paper [Tro11]:

$$\tilde{Z}^{ij} \equiv \begin{pmatrix} 0 & Z^{ij} \\ (Z^{ij})^T & 0 \end{pmatrix}.$$

Now we can apply the above lemma to these self-adjoint, independent and zero-mean random matrices.

To find R and $\tilde{A}^{ij}$'s that satisfy the conditions of the lemma, first consider a set of matrices $\{A^{ij}\}_{i<j}$ such that $\tilde{Z}^{ij} = \Delta_{ij} A^{ij}$ and

$$A^{ij} \;=\; \begin{pmatrix} 0 & e_i e_j^T - e_j e_i^T \\ e_j e_i^T - e_i e_j^T & 0 \end{pmatrix} ,$$

if $(i,j) \in E$ and zero otherwise. In the following, we show that the condition on p-th moment is satisfied with $R = 1/\sqrt{k d_{\max}^2}$ and $(\tilde{A}^{ij})^2 = (1/(k d_{\max}^2))(A^{ij})^2$ such that

$$\mathbb{E}\big[(\tilde{Z}^{ij})^p\big] \;\preceq\; \frac{p!}{2}\Big(\frac{1}{\sqrt{k d_{\max}^2}}\Big)^{p-2} \frac{1}{k d_{\max}^2} (A^{ij})^2 \,. \tag{28}$$

We can also show that $\tilde{\sigma}^2 \equiv \| \sum_{i<j} (\tilde{A}^{ij})^2 \|_2 = 1/(k d_{\max})$, since

$$\sum_{i<j} (\tilde{A}^{ij})^2 \;=\; \sum_{i<j} \frac{1}{k d_{\max}^2} \mathbb{I}_{((i,j)\in E)} \begin{pmatrix} e_i e_i^T + e_j e_j^T & 0 \\ 0 & e_i e_i^T + e_j e_j^T \end{pmatrix} \;=\; \frac{1}{k d_{\max}^2} \sum_{i=1}^n d_i \begin{pmatrix} e_i e_i^T & 0 \\ 0 & e_i e_i^T \end{pmatrix} ,$$

where $\mathbb{I}_{(\cdot)}$ is the indicator function. Therefore, we can apply the results of Lemma 4.4 to obtain a bound on $\big\| \sum_{i<j} Z^{ij} \big\|_2 = \big\| \sum_{i<j} \tilde{Z}^{ij} \big\|_2$:

$$\mathbb{P}\left(\left\| \sum_{i<j} Z^{ij} \right\| \geq t \right) \;\leq\; 2n \exp\left(\frac{-t^2/2}{(1/k d_{\max}) + (t/\sqrt{k d_{\max}^2})} \right) .$$

Under our assumption that $d_{\max} \geq \log n$ and choosing $t = C\sqrt{\log n/(k d_{\max})}$, the tail probability is bounded by $2n \exp\{-(C^2 \log n/2)(1/(1+C))\}$. Hence, we get the desired bound that $\|\Delta - D\|_2 \leq C\sqrt{\log n/(k d_{\max})}$ with probability at least $1 - 2n^{-(C^2-2C-2)/(2C+2)}$.

Now we are left to prove that the condition (28) holds. A quick calculation shows that

$$(A^{ij})^p = \begin{cases} (A^{ij})^2 & \text{for } p \text{ even} \,, \\ A^{ij} & \text{for } p \text{ odd} \,. \end{cases} \tag{29}$$

Furthermore, we can verify that the eigenvalues of A^{ij} are either 1 or -1. Hence, $(A^{ij})^p \preceq (A^{ij})^2$ for all $p \geq 1$. Thus, given the fact that $\tilde{Z}^{ij} = \Delta_{ij} A^{ij}$ we have that $\mathbb{E}[(\tilde{Z}^{ij})^p] = \mathbb{E}[\Delta_{ij}^p (A^{ij})^p] \preceq |\mathbb{E}[\Delta_{ij}^p]| (A^{ij})^2$ for all p. This fact follows since for any constant $c \in \mathbb{R}$, $cA^{ij} \preceq |c|(A^{ij})^2$ and $c(A^{ij})^2 \preceq |c|(A^{ij})^2$. Hence, coupling these observation with the identities presented in equation (29) we have

$$\mathbb{E}\big[(\tilde{Z}^{ij})^p\big] \;\preceq\; \mathbb{E}\big[|\Delta_{ij}|^p\big] (A^{ij})^2 \,,$$

where we used Jensen's inequality for $|\mathbb{E}[\Delta_{ij}^p]| \leq \mathbb{E}[|\Delta_{ij}|^p]$.

Next, it remains to construct a bound on $\mathbb{E}|\Delta_{ij}^p|$:

$$\mathbb{E}\big[|\Delta_{ij}|^p\big] \;\leq\; \frac{p!}{2}\Big(\frac{1}{\sqrt{k d_{\max}^2}}\Big)^p \,. \tag{30}$$

From (10), we have $\Delta_{ij} = P_{ij} - \tilde{P}_{ij} = \frac{1}{kd_{\max}} C_{ij}$. Therefore,

$$\mathbb{E}\big[|\Delta_{ij}|^p\big] = (1/kd_{\max})\mathbb{E}\big[|C_{ij}|^p\big].$$

Applying Azuma-Hoeffding's inequality to C_{ij}, we have that

$$\mathbb{P}\left(\frac{1}{kd_{\max}}|C_{ij}| \geq t\right) \leq 2\exp(-2t^2 d_{\max}^2 k)\ .$$

That is, $\frac{1}{kd_{\max}} C_{ij}$ is a sub-Gaussian random variable. And therefore, it follows that for $p \geq 2$,

$$\mathbb{E}\Big[\big|\frac{1}{kd_{\max}} C_{ij}\big|^p\Big] \ \leq \ \frac{p!}{2}\Big(\frac{1}{\sqrt{kd_{\max}^2}}\Big)^p\ .$$

This proves the desired bound in (30).

4.1.3 Proof of Lemma 4.3

By Lemma 4.2, we have

$$\begin{aligned} 1-\rho &= 1-\lambda_{\max}(\tilde{P}) - \|\Delta\|_2\sqrt{b} \\ &\geq 1-\lambda_{\max}(\tilde{P}) - C\sqrt{b\log n/(kd_{\max})}\ . \end{aligned}$$

In this section we prove that there is a positive gap: $(d_{\min}/2\,b^2\,d_{\max})\,\xi$. We will first prove that

$$1-\lambda_{\max}(\tilde{P}) \ \geq \ \frac{\xi\, d_{\min}}{b^2\, d_{\max}}\ . \tag{31}$$

This implies that we have the desired eigangap for $k \geq 4C^2b^5 d_{\max}\log n\,(1/d_{\min}\xi)^2$ such that $C\sqrt{b\log n/(kd_{\max})} \leq (d_{\min}/2\,b^2\,d_{\max})\,\xi$.

To prove (31), we use comparison theorems [DSC93], which bound the spectral gap of the Markov chain $\tilde{P}$ of interest using a few comparison inequalities related to a more tractable Markov chain, which is the simple random walk on the graph. We define the transition matrix of the simple random walk on the graph G as

$$Q_{ij} \ = \ \frac{1}{d_i} \quad \text{for } (i,j) \in E\ ,$$

and the stationary distribution of this Markov chain is $\mu(i) = d_i/\sum_j d_j$. Further, since the detailed balance equation is satisfied, Q is a reversible Markov chain. Formally, $\mu(i)Q_{ij} = 1/\sum_\ell d_\ell = \mu(j)Q_{ji}$ for all $(i,j) \in E$.

The following key lemma is a special case of a more general result [DSC93] proved for two arbitrary reversible Markov chains, which are not necessarily defined on the same graph. For completeness, we provide a proof of this lemma later in this section, following a technique similar to the one in [BGPS05] used to prove a similar result for a special case when the stationary distribution is uniform.

Lemma 4.5. *Let Q, μ and $\tilde{P}, \tilde{\pi}$ be reversible Markov chains on a finite set $[n]$ representing random walks on a graph $G = ([n], E)$, i.e. $\tilde{P}(i,j) = 0$ and $Q(i,j) = 0$ if $(i,j) \notin E$. For $\alpha \equiv \min_{(i,j)\in E}\{\tilde{\pi}(i)\tilde{P}_{ij}/\mu(i)Q_{ij}\}$ and $\beta \equiv \max_i\{\tilde{\pi}(i)/\mu(i)\}$,*

$$\frac{1-\lambda_{\max}(\tilde{P})}{1-\lambda_{\max}(Q)} \;\geq\; \frac{\alpha}{\beta}\,. \tag{32}$$

By assumption, we have $\xi \equiv 1 - \lambda_{\max}(Q)$. To prove that there is a positive spectral gap for the random walk of interest as in (31), we are left to bound α and β. We have $\mu(i)Q_{ij} = 1/\sum_\ell d_\ell \leq 1/|E|$ and $\mu(i) \geq (d_i/|E|)$. Also, by assumption that $\max_{i,j} w_i/w_j \leq b$, we have $\tilde{\pi}(i)\tilde{P}_{ij} = w_i w_j/(d_{\max}(w_i + w_j)\sum_\ell w_\ell) \geq 1/(bnd_{\max})$ and $\tilde{\pi}(i) = w_i/\sum_\ell w_\ell \leq b/n$. Then, $\alpha = \min_{(i,j)\in E}\{\tilde{\pi}(i)\tilde{P}_{ij}/\mu(i)Q_{ij}\} \geq |E|/(nbd_{\max})$ and $\beta = \max_i\{\tilde{\pi}(i)/\mu(i)\} \leq b|E|/nd_{\min}$. Hence, $\alpha/\beta \geq d_{\min}/(d_{\max}b^2)$ and this finishes the proof of the bound in (31).

4.1.4 Proof of Lemma 4.5

Since $1 - \lambda_{\max} = \min\{1 - \lambda_2, 1 + \lambda_n\}$, we will first show that $1 - \lambda_2(Q) \leq (\beta/\alpha)(1 - \lambda_2(\tilde{P}))$ and $1 + \lambda_n(Q) \leq (\beta/\alpha)(1 + \lambda_n(\tilde{P}))$. The desired bound in (32) immediately follows from the fact that $\min\{a, b\} \leq \min\{a', b'\}$ if $a \leq b$ and $a' \leq b'$.

A reversible Markov chain Q is self-adjoint in $L_2(\mu)$. Then, the second largest eigenvalue $\lambda_2(Q)$ can be represented by the Dirichlet form $\mathcal{E}$ defined as

$$\mathcal{E}^Q(\phi,\phi) \;\equiv\; \big\langle (I-Q)\phi, \phi\big\rangle_\mu \;=\; \frac{1}{2}\sum_{i,j}(\phi(i)-\phi(j))^2\mu(i)Q(i,j)\,.$$

For $\lambda_n(Q)$, we use

$$\mathcal{F}^Q(\phi,\phi) \;\equiv\; \big\langle (I+Q)\phi, \phi\big\rangle_\mu \;=\; \frac{1}{2}\sum_{i,j}(\phi(i)+\phi(j))^2\mu(i)Q(i,j)\,.$$

Following the usual variational characterization of the eigenvalues (see, for instance, [HJ85], p. 176) gives

$$1-\lambda_2(Q) \;=\; \min_{\phi\perp\mathbb{1}} \frac{\mathcal{E}^Q(\phi,\phi)}{\langle\phi,\phi\rangle_\mu}\,, \tag{33}$$

$$1+\lambda_n(Q) \;=\; \min_{\phi} \frac{\mathcal{F}^Q(\phi,\phi)}{\langle\phi,\phi\rangle_\mu}\,. \tag{34}$$

By the definitions of α and β, we have $\tilde{\pi}(i,)\tilde{P}(i,j) \geq \alpha\mu(i)Q(i,j)$ and $\tilde{\pi}(i) \leq \beta\mu(i)$ for all i and j, which implies

$$\begin{aligned}
\mathcal{E}^{\tilde{P}}(\phi,\phi) &\geq \alpha\mathcal{E}^Q(\phi,\phi)\,,\\
\mathcal{F}^{\tilde{P}}(\phi,\phi) &\geq \alpha\mathcal{F}^Q(\phi,\phi)\,,\\
\langle\phi,\phi\rangle_{\tilde{\pi}} &\leq \beta\langle\phi,\phi\rangle_\mu\,.
\end{aligned}$$

Together with (33), this implies $1 - \lambda_2(Q) \leq (\beta/\alpha)(1 - \lambda_2(\tilde{P}))$ and $1 + \lambda_n(Q) \leq (\beta/\alpha)(1 + \lambda_n(\tilde{P}))$. This finishes the proof of the desired bound.

4.2 Algorithm convergence for random sampling

Given the proof of Theorem 1 in the previous section, we only need to prove that for an Erdös-Renyi graph with average degree $d \geq C' \log n$ the following are true:

$$(1/2)d \;\leq\; d_i \;\leq\; (3/2)d\,, \tag{35}$$

$$1/2 \;\leq\; \xi\,. \tag{36}$$

Then, it follows that $\kappa \leq 3$ and $(1/2)d \leq d_{\max} \leq (3/2)d$. By Theorem 1, it follows that

$$\frac{\|\pi - \tilde{\pi}\|}{\|\tilde{\pi}\|} \;\leq\; Cb^{5/2}\sqrt{\frac{\log n}{k\,d}}\,,$$

for some positive constant C and for $kd \geq 288C^2b^5 \log n$ with probability at least $1 - n^{-C''}$.

We can apply standard concentration inequalities to establish equation (35). Apply Chernoff's inequality, we get $\mathbb{P}\big(\,|d_i - d| > (1/2)d\,\big) \leq 2e^{-d/16}$. Hence, for $d \geq C' \log n$, equation (35) is true with probability at least $1 - 2n^{-C'/16}$.

Finally, we finish the proof with a lower bound on the spectral gap $\xi = 1 - \lambda_{\max}(D^{-1}B)$. To establish this, we use celebrated results on the spectral gap of random graphs, first proved by Kahn and Szmerédie for d-regular random graphs [FKS89], and later extended to Erdös-Rényi graphs in [FO05]. Let B be the adjacency matrix of an Erdös-Rényi random graph $G(n, d/n)$, then it is shown in [FO05] that for $d \geq C' \log n$,

$$\sigma_2(B) \leq C\sqrt{d}\,, \tag{37}$$

with probability at least $1 - n^{-C''}$, for some numerical constants C, C', C'', where C'' can be made as large as we want by increasing C and C'. Since we are interested in the eigenvalues of $L = D^{-1}B$, we define a more tractable matrix with the same set of eigenvalues: $\tilde{L} = D^{-1/2}BD^{-1/2}$. Because $\tilde{L}$ is a symmetric matrix, the eigenvalues are the same as the singular values up to a sign. Let $\sigma_1(\tilde{L}) \geq \sigma_2(\tilde{L}) \geq \ldots$ denote the ordered singular values of $\tilde{L}$. Then, it is enough to show that

$$\lambda_{\max}(L) \;=\; \sigma_2(\tilde{L}) \;\leq\; \frac{2}{d}\sigma_2(B)\,. \tag{38}$$

Then,

$$\xi \;=\; 1 - \lambda_{\max}(L) \;\geq\; 1 - \frac{2C}{\sqrt{d}}\,,$$

and for $d \geq C' \log n$, this can be made as close to one as we want by increasing C'. This proves the desired lower bound in (36).

We are left to prove the bound in (38) using the variational representation of $\sigma_2(\tilde{L})$:

$$\begin{aligned}
\sigma_2(\tilde{L}) &= \min_{H \in \mathcal{H}_2} \max_{x \in H} \frac{x^T D^{-1/2} B D^{-1/2} x}{x^T x} \\
&= \min_{H \in \mathcal{H}_2} \max_{y \in H} \frac{y^T B y}{y^T D y} \\
&\leq \min_{H \in \mathcal{H}_2} \max_{y \in H} \frac{1}{d_{\min}} \frac{y^T B y}{y^T y} \\
&= \frac{\sigma_2(B)}{d_{\min}}\,,
\end{aligned}$$

where $\mathcal{H}_2$ is the set of all 2-dimensional subspaces in $\mathbb{R}^n$ and $d_{\min}$ is the minimum degree. By concentration of measure inequalities, we know that $d_{\min} \geq (1/2)d$.

4.3 Proof of the information-theoretic lower bound in Theorem 3

In this section, we prove Theorem 3 using an information-theoretic method that allows us to reduce the stochastic inference problem into a multi-way hypothesis testing problem.

This estimation problem can be reduced to the following hypothesis testing problem. Consider a set $\{\tilde{\pi}^{(1)}, \ldots, \tilde{\pi}^{(M(\delta))}\}$ of $M(\delta)$ vectors on the standard orthogonal simplex which are separated by δ, such that $\|\tilde{\pi}^{(\ell_1)} - \tilde{\pi}^{(\ell_2)}\| \geq \delta$ for all $\ell_1 \neq \ell_2$. To simplify the notations, we are going to use M as a shorthand for $M(\delta)$. Suppose we choose an index $L \in \{1, \ldots, M\}$ uniformly at random. Then, we are given noisy outcomes of pairwise comparisons with $w = \tilde{\pi}^{(L)}$ from the BTL model. We use X to denote this set of observations. Let π be the estimation produced by an algorithm using the noisy observations. Given this, the best estimation of the "index" is $\hat{L}$, where $\hat{L} = \arg\min_{\ell \in [M]} \|\pi - \tilde{\pi}^{(\ell)}\|$.

By construction of our packing set, when we make a mistake in the hypothesis testing, our estimate is at least $\delta/2$ away from the true weight $\tilde{\pi}^{(L)}$. Precisely, $\hat{L} \neq L$ implies that $\|\pi - \tilde{\pi}^{(L)}\| \geq \delta/2$. Then,

$$\begin{aligned} \mathbb{E}\big[\, \|\pi - \tilde{\pi}^{(L)}\| \,\big] &\geq \frac{\delta}{2}\mathbb{P}\big(\hat{L} \neq L\big) \\ &\geq \frac{\delta}{2}\Big\{1 - \frac{I(\hat{L}; L) + \log 2}{\log M}\Big\}\,, \end{aligned} \tag{39}$$

where $I(\cdot;\cdot)$ denotes the mutual information between two random variables and the second inequality follows from Fano's inequality.

These random vectors form a Markov chain $L \multimap \tilde{\pi}^{(L)} \multimap X \multimap \pi \multimap \hat{L}$, where $X \multimap Y \multimap Z$ indicates that X and Z are conditionally independent given Y. Let $\mathbb{P}_{L,X}(\ell, x)$ denote the joint probability function, and $\mathbb{P}_{X|L}(x|\ell)$, $\mathbb{P}_L(\ell)$ and $\mathbb{P}_X(x)$ denote the conditional and marginal probability functions. Then, by data processing inequality for a Markov chain, we get

$$\begin{aligned} I(L; \hat{L}) &\leq I(L; X) \\ &= \mathbb{E}_{L,X}\Big[\log\Big(\frac{\mathbb{P}_{L,X}(L, X)}{\mathbb{P}_L(L)\mathbb{P}_X(X)}\Big)\Big] \\ &= \frac{1}{M}\sum_{\ell \in [M]} \mathbb{E}_X\Big[\log\Big(\frac{\mathbb{P}_{X|L}(X|\ell)}{\mathbb{P}_X(X)}\Big)\Big] \\ &\leq \frac{1}{M^2}\sum_{\ell_1, \ell_2} D_{\mathrm{KL}}\Big(\, \mathbb{P}_{X|L}(X|\ell_1) \,\Big\|\, \mathbb{P}_{X|L}(X|\ell_2) \,\Big)\,, \end{aligned} \tag{40}$$

where $D_{\mathrm{KL}}(\cdot\|\cdot)$ is the Kullback-Leibler (KL) divergence and the last inequality follows from the convexity of KL divergence and Jensen's inequality.

The KL divergence between the observations coming from two different BTL models depend on how we sample the comparisons. We are sampling each pair of items for comparison with probability d/n, and we are comparing each of these sampled pairs k times. Let X_{ij} denote the outcome of k comparisons for a sampled pair of items (i, j). To simplify notations, we drop the

subscript $X|L$ whenever it is clear from the context. Then,

$$\begin{aligned} D_{\mathrm{KL}}\big(\,\mathbb{P}(X|\ell_1)\,\big\|\,\mathbb{P}(X|\ell_2)\,\big) &= \frac{d}{n}\sum_{1\le i<j\le n} D_{\mathrm{KL}}\big(\mathbb{P}(X_{ij}|\ell_1)\,\big\|\,\mathbb{P}(X_{ij}|\ell_2)\big) \\ &\le 2\,n^2\,k\,d\,\big\|(\tilde\pi^{(\ell_1)}-\tilde\pi^{(\ell_2)}\big\|^2\,, \end{aligned} \tag{41}$$

where in the last inequality we used the fact that

$$\begin{aligned} D_{\mathrm{KL}}\big(\,\mathbb{P}(X_{ij}|\ell_1)\,\|\,\mathbb{P}(X_{ij}|\ell_2)\,\big) &\le \frac{k\big(\tilde\pi_j^{(\ell_2)}(\tilde\pi_i^{(\ell_1)}-\tilde\pi_i^{(\ell_2)})^2+\tilde\pi_i^{(\ell_2)}(\tilde\pi_j^{(\ell_1)}-\tilde\pi_j^{(\ell_2)})^2\big)}{\tilde\pi_i^{(\ell_2)}\tilde\pi_j^{(\ell_2)}\big(\tilde\pi_i^{(\ell_1)}+\tilde\pi_j^{(\ell_1)}\big)} \\ &\le 2kn^2\Big((\tilde\pi_i^{(\ell_1)}-\tilde\pi_i^{(\ell_2)})^2+(\tilde\pi_j^{(\ell_1)}-\tilde\pi_j^{(\ell_2)})^2\Big)\,, \end{aligned}$$

for k independent trials of Bernoulli random variables, and $\tilde\pi_i^{(\ell)} \ge 1/(2n)$ for all i and ℓ which follows from our construction of the packing set in Lemma 4.6 and our choice of δ.

The remainder of the proof relies on the following key technical lemma, on the construction of a suitable packing set that has enough number of entries which are reasonably separated. This is proved in Section 4.3.1.

Lemma 4.6. *For $n \ge 90$ and for any positive $\delta \le 1/2\sqrt{10n}$, there exists a set of n-dimensional vectors $\{\tilde\pi^{(1)},\dots,\tilde\pi^{(M)}\}$ with cardinality $M = e^{n/128}$ such that $\sum_i \tilde\pi_i^{(\ell)} = 1$ and*

$$\frac{1-2\delta\sqrt{10n}}{n} \le \tilde\pi_i^{(\ell)} \le \frac{1+2\delta\sqrt{10n}}{n}\,,$$

for all $i \in [n]$ and $\ell \in [M]$, and

$$\delta \le \|\tilde\pi^{(\ell_1)}-\tilde\pi^{(\ell_2)}\| \le \sqrt{13}\delta\,,$$

for all $\ell_1 \ne \ell_2$.

Substituting this bound in Eqs. (41), (40), and (39), we get

$$\begin{aligned} \max_{\ell\in[M]} \mathbb{E}[\,\|\pi-\tilde\pi^{(\ell)}\|\,] &\ge \mathbb{E}[\,\|\pi-\tilde\pi^{(L)}\|\,] \\ &\ge \frac{\delta}{2}\Big\{1-\frac{3328n^2kd\delta^2+128\log 2}{n}\Big\}\,. \end{aligned}$$

Choosing $\delta = (b-1)/(30\sqrt{10}(b+1)\sqrt{kdn})$, we know that $3328n^2kd\delta^2 + 128\log 2 \le (1/2)n$ for all b and all $n \ge 682$. This implies that

$$\max_{\ell\in[M]} \mathbb{E}\big[\,\|\pi-\tilde\pi^{(\ell)}\|\,\big] \ge \frac{(b-1)}{120(b+1)\sqrt{10kdn}}\,.$$

From Lemma 4.6, it follows that $\|\tilde\pi^{(\ell)}\| \le 2/\sqrt{n}$ for all ℓ. Then, scaling the bound by $1/\|\tilde\pi^{(\ell)}\|$, the normalized minimax rate is lower bounded by $(b-1)/(240(b+1)\sqrt{10kd})$. Also, for this choice of δ, the dynamic range is at most b. From Lemma 4.6, the dynamic range is upper bounded by

$$\max_{\ell,i,j}\frac{\tilde\pi_i^{(\ell)}}{\tilde\pi_j^{(\ell)}} \le \frac{1+2\delta\sqrt{10n}}{1-2\delta\sqrt{10n}}\,.$$

This is monotonically increasing in δ for $\delta < 1/(2\sqrt{10n})$. Hence, for $\delta \le (b-1)/((b+1)2\sqrt{10n})$, which is always true for our choice of δ, the dynamic range is upper bounded by b. This finishes the proof of the desired bound on normalized minimax error rate for general b.

4.3.1 Proof of Lemma 4.6

We show that a random construction succeeds in generating a set of M vectors on the standard orthogonal simplex satisfying the conditions with a strictly positive probability. Let $M = e^{n/128}$ and for each $\ell \in [M]$, we construct independent random vectors $\tilde{\pi}^{(\ell)}$ according to the following procedure. For a positive α to be specified later, we first draw n random variables uniformly from $[(1 - \alpha\delta\sqrt{n})/n, (1 + \alpha\delta\sqrt{n}))/n]$. Let $Y^{(\ell)} = [Y_1^{(\ell)}, \ldots, Y_n^{(\ell)}]$ denote this random vector in n dimensions. Then we project this onto the n-dimensional simplex by setting

$$\tilde{\pi}^{(\ell)} \;=\; Y^{(\ell)} + (1/n - \bar{Y}^{(\ell)})\mathbb{1} \;,$$

where $\bar{Y}^{(\ell)} = (1/n)\sum_i Y_i^{(\ell)}$. By construction, the resulting vector is on the standard orthogonal simplex: $\sum_i \tilde{\pi}_i^{(\ell)} = 1$. Also, applying Hoeffding's inequality for $\bar{Y}^{(\ell)}$, we get that

$$\mathbb{P}\Big(\Big|\bar{Y}^{(\ell)} - \frac{1}{n}\Big| > \frac{\alpha\delta}{\sqrt{n}}\Big) \;\leq\; 2e^{-n/2} \;.$$

By union bound, this holds uniformly for all ℓ with probability at least $1 - 2e^{-63n/128}$. In particular, this implies that

$$\frac{1 - 2\alpha\delta\sqrt{n}}{n} \;\leq\; \tilde{\pi}_i^{(\ell)} \;\leq\; \frac{1 + 2\alpha\delta\sqrt{n}}{n} \;, \tag{42}$$

for all $i \in [n]$ and $\ell \in [M]$.

Next, we use standard concentration results to bound the distance between two vectors:

$$\big\|\tilde{\pi}^{(\ell_1)} - \tilde{\pi}^{(\ell_2)}\big\|^2 \;=\; \big\|Y^{(\ell_1)} - Y^{(\ell_2)}\big\|^2 - n(\bar{Y}^{(\ell_1)} - \bar{Y}^{(\ell_2)})^2$$

Applying Hoeffding's inequality for the first term, we get $\mathbb{P}\big(|\sum_i (Y_i^{(\ell_1)} - Y_i^{(\ell_2)})^2 - (2/3)\alpha^2\delta^2| \geq (1/2)\alpha^2\delta^2\big) \leq 2e^{-n/32}$. Similarly for the second term, we can show that $\mathbb{P}\big(|\sum_i (Y_i^{(\ell_1)} - Y_i^{(\ell_2)})| \geq (1/4)\alpha\delta\sqrt{n}\big) \leq 2e^{-n/32}$. Substituting these bounds, we get

$$\frac{1}{10}\alpha^2\delta^2 \;\leq\; \|\tilde{\pi}^{(\ell_1)} - \tilde{\pi}^{(\ell_2)}\|^2 \;\leq\; \frac{13}{10}\alpha^2\delta^2 \;, \tag{43}$$

with probability at least $1 - 4e^{-n/32}$. Applying union bound over $\binom{M}{2} \leq e^{n/64}$ pairs of vectors, we get that the lower and upper bound holds for all pairs $\ell_1 \neq \ell_2$ with probability at least $1 - 4e^{-n/64}$.

The probability that both conditions (42) and (43) are satisfied is at least $1 - 4e^{-n/64} - 2e^{-63n/128}$. For $n \geq 90$, the probability of success is strictly positive. Hence, we know that there exists at least one set of vectors that satisfy the conditions. Setting $\alpha = \sqrt{10}$, we have constructed a set that satisfy all the conditions.

4.4 Proof of Lemma 3.1

Without loss of generality, let us consider two items i and j such that $w_i > w_j$. When we estimate a higher score for item j then we make a mistake in the ranking of these two items. When this happens, such that $\pi_j - \pi_i > 0$, it naturally follows that $w_i - w_j \leq w_i - w_j + \pi_j - \pi_i \leq |w_i - \pi_i| + |\pi_j - w_j|$. For a general pair i and j, we have $(w_i - w_j)(\sigma_i - \sigma_j) > 0$ implies that

$|w_i - w_j| \leq |w_i - \pi_i| + |w_j - \pi_j|$. Substituting this into the definition of the weighted distance $D_w(\cdot)$, and using the fact that $(a+b)^2 \leq 2a^2 + 2b^2$, we get

$$\begin{aligned} D_w(\sigma) &= \Big\{ \frac{1}{2n\|w\|^2} \sum_{i<j} (w_i - w_j)^2 \, \mathbb{I}\big((w_i - w_j)(\sigma_i - \sigma_j) > 0\big) \Big\}^{1/2} \\ &\leq \Big\{ \frac{1}{n\|w\|^2} \sum_{i<j} \big\{ (w_i - \pi_i)^2 + (w_j - \pi_j)^2 \big\} \Big\}^{1/2} \\ &\leq \frac{1}{\|w\|} \sum_{i=1}^{n} (w_i - \pi_i)^2 \, . \end{aligned}$$

This proves that the distance $D_w(\sigma)$ is upper bounded by the normalized Euclidean distance $\|w - \pi\|/\|w\|$.

5 Discussion

In this paper, we developed a novel iterative rank aggregation algorithm for discovering scores of objects given pairwise comparisons. The algorithm has a natural *random walk* interpretation over the graph of objects with edges present between two objects if they are compared; the scores turn out to be the stationary probability of this random walk. In lieu of recent works on network centrality which are graph score functions primarily based on random walks, we call this algorithm *Rank Centrality*. The algorithm is model independent.

We also established the efficacy of the algorithm by analyzing its performance when data is generated as per the popular Bradley-Terry-Luce (BTL) model. We have obtained an analytic bound on the finite sample error rates between the scores assumed by the BTL model and those estimated by our algorithm. As shown, these lead to order-optimal dependence on the number of samples required to learn the scores well by our algorithm under random select of pairs for comparison. The experimental evaluation show that our (model independent) algorithm performs as well as the Maximum Likelihood Estimator of the BTL model and outperforms other known competitors included the recently proposed algorithm by Ammar and Shah [AS11].

Given the simplicity of the algorithm, analytic guarantees and wide utility of the problem of rank aggregation, we strongly believe that this algorithm will be of great practical value.

References

[AGHB+94] M. Adler, P. Gemmell, M. Harchol-Balter, R. M. Karp, and C. Kenyon. Selection in the presence of noise: the design of playoff systems. In *Proceedings of the fifth annual ACM-SIAM symposium on Discrete algorithms*, SODA '94, pages 564–572. Society for Industrial and Applied Mathematics, 1994.

[Arr63] K. J Arrow. *Social Choice and Individual Values*. Yale University Press, 1963.

[AS11] A. Ammar and D. Shah. Ranking: Compare, don't score. In *Communication, Control, and Computing (Allerton), 2011 49th Annual Allerton Conference on*, pages 776–783, September 2011.

[BGPS05] S. Boyd, A. Ghosh, B. Prabhakar, and D. Shah. Mixing times for random walks on geometric random graphs. *SIAM ANALCO*, 2005.

[BM08] M. Braverman and E. Mossel. Noisy sorting without resampling. In *Proceedings of the nineteenth annual ACM-SIAM symposium on Discrete algorithms*, SODA '08, pages 268–276. Society for Industrial and Applied Mathematics, 2008.

[BP98] S. Brin and L. Page. The anatomy of a large-scale hypertextual web search engine. In *Seventh International World-Wide Web Conference (WWW 1998)*, 1998.

[BT55] R. A. Bradley and M. E. Terry. Rank analysis of incomplete block designs: I. the method of paired comparisons. *Biometrika*, 39(3/4):324–345, 1955.

[Con85] M. Condorcet. *Essai sur l'application de l'analyse à la probabilité des décisions rendues à la pluralité des voix*. l'Imprimerie Royale, 1785.

[CR09] E. J. Candès and B. Recht. Exact matrix completion via convex optimization. *Foundations of Computational Mathematics*, 9(6):717–772, 2009.

[DKNS01] C. Dwork, R. Kumar, M. Naor, and D. Sivakumar. Rank aggregation methods for the web. In *Proceedings of the Tenth International World Wide Web Conference, 2001*, 2001.

[DMJ10] J. C. Duchi, L. Mackey, and M. I. Jordan. On the consistency of ranking algorithms. In *Proceedings of the ICML Conference*, Haifa, Israel, June 2010.

[DSC93] P. Diaconis and L. Saloff-Coste. Comparison theorems for reversible markov chains. *The Annals of Applied Probability*, 3(3):696–730, 1993.

[FKS89] J. Friedman, J. Kahn, and E. Szemerédi. On the second eigenvalue in random regular graphs. In *Proceedings of the Twenty-First Annual ACM Symposium on Theory of Computing*, pages 587–598, Seattle, Washington, USA, may 1989. ACM.

[FO05] U. Feige and E. Ofek. Spectral techniques applied to sparse random graphs. *Random Struct. Algorithms*, 27(2):251–275, 2005.

[HJ85] R. A. Horn and C. R. Johnson. *Matrix Analysis*. Cambridge University Press, 1985.

[KMO10] R. H. Keshavan, A. Montanari, and S. Oh. Matrix completion from noisy entries. *Journal of Machine Learning Research*, 11:2057–2078, July 2010.

[KSGM03] S. D. Kamvar, M. T. Schlosser, and H. Garcia-Molina. The eigentrust algorithm for reputation management in p2p networks. In *Proceedings of the 12th international conference on World Wide Web*, WWW '03, pages 640–651, New York, NY, USA, 2003. ACM.

[LRF57] Jr. L. R. Ford. Solution of a ranking problem from binary comparisons. *The American Mathematical Monthly*, 64(8):28–33, 1957.

[Luc59] D. R. Luce. *Individual Choice Behavior*. Wiley, New York, 1959.

[McF73] D. McFadden. Conditional logit analysis of qualitative choice behavior. *Frontiers in Econometrics*, pages 105–142, 1973.

[New10] M. E. J. Newman. *Networks: An Introduction*. Oxford University Press, 2010.

[NW12] S. Negahban and M. J. Wainwright. Restricted strong convexity and (weighted) matrix completion: Optimal bounds with noise. *Journal of Machine Learning Research*, pages 1665–1697, May 2012.

[Saa03] T. L. Saaty. Decision-making with the ahp: Why is the principal eigenvector necessary. *European Journal of Operational Research*, 145:pp. 85–91, 2003.

[SL12] M. J. Salganik and K. E.C. Levy. Wiki surveys: Open and quantifiable social data collection. Technical Report `arXiv:1202.0500`, 2012.

[TR05] K. T. Talluri and G. Van Ryzin. *The Theory and Practice of Revenue Management*. springer, 2005.

[Tro11] J. Tropp. User-friendly tail bounds for sums of random matrices. *Foundations of Computational Mathematics*, 2011.